Strategic Plan for Grids (2010-2012)

Eileen Berman, Keith Chadwick, Stu Fuess, Gabriele Garzoglio, Burt Holzman, Don Petravick, Ruth Pordes, Stephen Wolbers

V1.1, 08/13/09

Fermilab's scientific program involves collaborations that are national or international in scale. Computing for these activities is spread among many facilities[1], and a single facility needs to serve many experiments, projects, and collaborations.

Grid (distributed) computing seeks to organize computing along the roles of these experiments, projects, and collaborations, ("virtual organizations"), and resource providers ("sites"[2]) One of the main goals of grid computing is to reduce the coupling between sites and virtual organizations to a minimal number of transparent, standard interfaces and business processes.

Virtual organizations gain from this approach in that they are able to add sites at low marginal cost, and sustain operations. Fermilab gains because it is able to support computing for new virtual organizations at low marginal cost. Both parties gain because the coupling between sites and virtual organizations is reduced to a thin interface, allowing each party to focus on optimizing its processes with relative independence.

Grid techniques are required by the Large Hadron Collider (LHC) software community[3], therefore embracing them is a necessity. Improving upon the current set of techniques and identifying tractable additions and expansions would provide benefit for Fermilab virtual organizations and our external grid collaborators. Primary reference points for locating such mechanisms are the Fermilab Computing Division and Fermilab-hosted experiments/projects, the Open Science Grid (OSG) and the Worldwide LHC Computing Grid (WLCG.) Constraints on acceptable techniques may come from cyber security.

The grid approach creates maximal value when it is broadly accepted. It is necessary to both collaborate with others and to build on our existing strengths. This will both help the virtual organizations and strengthen our ties into a national cyberinfrastructure, enabling us to keep pace with this infrastructure and participate in its evolution.

Fermilab has a history of working closely with experiments/projects, thinking carefully about operations, interoperable (and especially default-allow) security, data movement and storage. These skills will contribute positively to a grid effort which must apply the knowledge gained from the acquisition of these skills to creation and operation of the grid infrastructure itself.

1 A set of resources that provide a particular service, e.g. grid facility, storage facility.
2 An infrastructure of heterogeneous resources under a single administrative domain. A single site may contain multiple facilities.
3 CMS, ATLAS, ALICE, LHCb

The Fermilab Computing Division faces the challenge of supporting the Run II experiments in a period of increasing luminosity and decreasing available effort, assisting CMS during the critical first years at the new energy frontier, facilitating the use of core grid services of the future programs of Fermilab and maintaining flexibility to understand and participate in the evolution of distributed computing in support of the Fermilab scientific program.

Mission

To enhance and expand the body of grid software, business methods, and deployment community that is broadly accepted by the Fermilab site and Fermilab based virtual organizations.

To expand and improve upon the relationship that exists between Fermilab based virtual organizations, distributed computing infrastructures (grids) and sites.

To assess and measure the acceptance and effectiveness of these techniques.

To perform targeted research and development to enhance the methods and technologies used in Grid production systems, specifically including computer security, in our distributed open science environment.[4]

Assessment of Current State (2009)

Fermilab offers a production grid facility (FermiGrid), which provides a set of common services and a portal to the OSG. Through this facility, virtual organizations access both FermiGrid and OSG resources. CDF, CMS and DØ generate large MC samples, CDF runs analysis jobs on FermiGrid and OSG resources and DØ has reconstructed more than 1B events and runs primary production on these resources. FermiGrid and OSG are utilized by many experiments/projects at Fermilab. Stakeholder interoperability and opportunistic use of resources commonly occur. Support of OSG virtual organizations is ongoing. FermiGrid is one of the preeminent computing facilities in the OSG. OSG communities have access to significant opportunistic computing resources at Fermilab (~150,000 CPU hours/month). Operational security practices and procedures governing security at the grid-site boundary are being investigated and realized. The continued deployment of high availability techniques is contributing to the strengthening of the distributed infrastructure and increased availability of the offered services. Extensions to the deployed automated monitoring have increased timely response to operational incidents. The SAZ[5] and GUMS[6] supported service levels were increased to 24x7

[4] From Wikipedia: Open research is research, conducted in the spirit of free and open source software. If the research is scientific in nature it is frequently referred to as open science; open research can also include social sciences and the humanities.
[5] Site AuthoriZation Service
[6] Grid User Management Service

support. Application of ITIL (Information Technology Infrastructure Library) techniques is continuing and will contribute to the understanding and operation of the infrastructure as an enterprise level activity.

CD is actively involved in grid middleware development, integration, deployment, and support for authorization and identity management, accounting, resource and workload management, data handling, and storage and data movement. Many of these activities involve collaboration with external Grid projects and organizations. The VO Services Project[7] and SAMGrid Project[8] have wound down development and are focused mainly on operations. CD expertise in the above areas is an often-requested capability. Acceptance and use of the glideinWMS[9] deliverables has expanded to additional Fermilab VOs. Awareness and application of good security practices and policies is increasing and contributing to worldwide efforts. Strengthening of the grid infrastructure to production quality has progressed and many components are being heavily used in production. Investigation of enterprise architecture techniques and their application to distributed computing middleware (including generation of metrics-based reports) has begun. Application of these techniques in the area of service monitoring and troubleshooting is being investigated. User support is ongoing and dedicated effort is available for problem triaging and user assistance.

Vision

To integrate HEP computing at US universities to the common grid distributed computing infrastructure. This allows coherent HEP computing to emerge in the US, and ensures that Fermilab is anchored in this infrastructure, allowing it to evolve its computing facilities and expertise in a way that gives maximal advantage to the US HEP community.

To influence the labs non-HEP program (QCD, Astrophysics) to evolve along similar or identical lines.

To be seamlessly integrated into the distributed computing system for CMS as a Tier-1 center and as an analysis facility including the production deployment of a workload management system. Processing resources for event reconstruction, data selection, and regional analysis will be transparently accessible through the OSG infrastructure and will have been scaled to support the additional resource needs. Custodial data storage from the experiment and event simulation centers, as well as data serving to remote analysis facilities will be accessible through reliable and secure common grid interfaces.

7 Provides software solutions for VO user registration and fine-grained authorization for access to grid-enabled resources

8 A general data handling system designed for experiments/projects with large (petabyte-sized) datasets and widely distributed production and analysis facilities

9 Condor glidein workload management system that provides a simple mechanism for accessing grid resources

In the coming few years, the Run II experiments will have completed their raw data taking and will be focusing on Monte Carlo generation and data analysis. Migration to stable operations will have occurred. This includes a set of core computing services that perform at the scale required to support these activities. Computing services will have been migrated to common grid solutions where this is seen as the most efficient program of work, weighing the development cost against the potential savings in operations and support.

For all experiments/projects where legacy solutions are operated, any transition to maintenance will include the identification of sufficient effort to provide adequate support for the program of the virtual organization.

Fermilab will provide, in addition to the current advanced set of core services, a production set of distributed computing workload management tools in use by multiple virtual organizations, an enterprise level grid aware metrics aggregation service, and monitoring and alarming infrastructure, and the deployment of High Performance Computing (HPC) capabilities for distributed computing. The work with the storage implementation in the Virtual Data Toolkit[10] will enable Fermilab to continue as the primary expert in grid enabled storage solutions for the OSG. Fermilab will have sufficient effort and expertise to facilitate and guide the transition of smaller communities to common grid services and components where desired.

Grid operation will have developed into a long-term sustainable activity with minimal operational effort. Available operational tools will provide information to operations teams, security personnel, and grid users. Virtualization technologies will be used to improve the user experience and capabilities and ease operational load. There will be seamless interoperation with alternate authentication mechanisms (such as Shibboleth, used in university campus grids). Commercial cloud technologies will be interfaced to, as appropriate, to enhance and extend users access to available resources.

Contributions will be made to research and development to enhance the capabilities and effectiveness of Grid computing in areas of computer security, production distributed infrastructures, and large-scale data, high throughput grids.

Stakeholders

The sponsors of the grid work are the Fermilab Computing Division base program and the DOE and NSF in the SciDAC-2 and NSF sponsored projects (OSG, CEDPS). Contributions are made from the US CMS software and computing project. Effort and deliverables are provided from many groups both internal and external to the Laboratory.

The stakeholders are:

- The running experiments based at Fermilab (CDF, D0, Minos, MiniBooNE), the neutrino program, and the astrophysics community.
- CMS.

10 VDT – The software stack containing the OSG grid software (services, clients, tools.)

- Simulation and theory including accelerator modeling, LQCD.
- The Computing Division.
- The members of the Open Science Grid Consortium.
- The Worldwide LHC Computing Grid (WLCG).
- The greater US HEP community, including QCD and astrophysics communities.
- The SBIR community.

Goals and Objectives

- Perform all work activities with awareness and understanding of appropriate safety practices and procedures.
- Operate, support, and evolve a robust, effective, secure, local production grid facility, which supports the scientific program of Fermilab (FermiGrid).
- Increase FermiGrid's operational efficiency by the application of High Availability and ITIL techniques, enhanced monitoring and alarming, and the creation of additional operational tools.
- Provide a grid services platform including services that are contributions to the OSG and help Fermilab achieve its goals. Special attention should be paid to implementations compatible with the infrastructure as outlined in the CMS Computing Technical Design Report (CTDR). Collaborate with external projects in the development and deployment of these grid services and grid tools.
- Understand and develop grid security policies, processes, and procedures in order to ensure the secure operation of the production grid facility. Assist the experiments/projects in developing operational security expertise.
- Understand and work to help provide grid accessible storage solutions that provide for the needs of Fermilab based VO’s.
- Develop security processes and procedures surrounding the software development life cycle. Apply these processes and procedures to existing and future projects.
- Work towards seamless interoperation between the national and regional grids that Run II and CMS worldwide collaborations depend on– in particular the EGEE (Enabling Grids for E-sciencE Project) and OSG.
- Work towards interoperation with additional national and regional grids such as TeraGrid and campus grids.
- Maintain visibility as a leading member of the OSG, used by all Fermilab scientific groups and take leadership roles within this consortium.
- Maintain and develop relationships and collaboration with sites, software contributors, committees and other bodies sufficient to increase the usability and administration of grid infrastructure integral to the Fermilab mission. Special attention should be paid to the needs of the DOE environment.
- Understand/Improve the relationship between VOs, resource providers, and Grid providers and the user experience working with grid tools.
- Offer excellent support to the Fermilab experiments/projects to enhance and enable their efficient use of FermiGrid, OSG, and worldwide grid facilities. Provide a support level that is appropriate to enable the virtual organization to operate in an efficient manner.

- Promote use of distributed computing facilities by additional experiments/projects including the neutrino and astrophysics programs as needed.
- Provide outreach and communication activities to further the movement of Fermilab experiments/projects to using the distributed grid infrastructure.
- Investigate the application of MPI technologies in the grid environment. Deploy these technologies in the grid environment to support the MPI applications.
- Investigate the utility of scientific cloud computing to the grid community and apply to distributed computing at Fermilab as appropriate.
- Investigate the utility of commercial cloud offerings to the grid community and utilize these offerings at Fermilab as appropriate.
- Investigate the utility of virtualization to the grid community and apply to distributed computing at Fermilab as appropriate.
- Identify a grid/distributed computing based path forward for HEP and astrophysics communities.
- Contribute to research and development collaborations for advancing computer security for open science.

Strategies

Work with specific experiments/projects to both gain domain knowledge and foster acceptance of the distributed computing approach.

Use CMS and the Run II experiments as exemplars and critical leverage to evolve and expand distributed computing techniques and technologies.

Undertake leadership and supporting roles in various organizations, including the Open Science Grid, and evolving the organization to build acumen in applicable techniques.

Partner and collaborate with other significant contributors in developing and deploying distributed computing middleware and production facility capabilities.

Provide common grid services at Fermilab, including the operation of the FermiGrid campus grid.

Evolve site operational strengths to a distributed computing context.

Work with other departments in CD to facilitate the deployment of grid software and its operation and support.

Collaborate with computer science, mathematics, and domain science groups in the DOE Office of Science and the NSF to advance the state of the art, improve the effectiveness and mitigate the risks associated with open science distributed computing

Resource Needs

All resources purchased for the LHC will be entirely consumed with LHC activities, so there should not be any planned reliance on securing large numbers of opportunistic computing resources from these resources.

FermiGrid resources can be expected to require growth over the next three years, both in the area of service machines (i.e. Gatekeepers) and worker nodes.

CMS effort will contribute to both the experiment specific services and to the common grid services and infrastructures, increasingly shifting to operations from development over the next 2 years.

With the increase in luminosity for Run II, and the increasing interest in workload management systems by more experiments/projects, support of these users and the resources they require will occupy a larger fraction of effort.

Increase in the types and variety of storage services necessary for efficient and effective grid use by virtual organizations and the associated use cases will require increased effort.

Progress Indicators

Development project milestones will be tracked and reported. For virtual organization layers, compliance with standards will be tracked. The following will be reported monthly, taking account of the virtual organization needs:
-- The % of jobs through grid interfaces per virtual organization.
-- The number of different types of jobs using grid interfaces.
-- Number of sites offering turnkey grid access per virtual organization.
-- Efficiency of use of grid resources.
-- Numbers of incidents a week and the successful resolution of said incidents.
-- Number of problems involving associated services.
-- Decreasing time spent in administration of grid services and resources.
-- Decreasing time spent by the users to get up to speed and sustain use of the services and resources.

Risks

The risks of this strategic plan include:

- The continued support and attention of the underlying Grid technology groups – Condor (batch system support for highly distributed computing) and Globus (toolkit for building grid systems) – to deliver the standard middleware used by all grids.
- Failure at the annual review for funding of the Open Science Grid.
- Failures in communication with the virtual organizations to ensure their ongoing requirements and plans are met and/or insufficient resources made available to meet their goals.

- Failure of internal communication within a virtual organization such that the Fermilab grid activities have no reliable communication mechanism with the VO.
- Failure to have adequate staffing levels to meet needs and/or goals.
- CD funding levels may be inadequate for the needed ongoing hardware replacement/upgrades to manage hardware lifecycles.